

Mobile Geometric Graphs, and Detection and Communication Problems in Mobile Wireless Networks

Alistair Sinclair*

Alexandre Stauffer†

July 8, 2010

Abstract

Static wireless networks are by now quite well understood mathematically through the *random geometric graph* model. By contrast, there are relatively few rigorous results on the practically important case of *mobile* networks, in which the nodes move over time; moreover, these results often make unrealistic assumptions about node mobility such as the ability to make very large jumps. In this paper we consider a realistic model for mobile wireless networks which we call *mobile geometric graphs*, and which is a natural extension of the random geometric graph model. We study two fundamental questions in this model: *detection* (the time until a given “target” point—which may be either fixed or moving—is detected by the network), and *percolation* (the time until a given node is able to communicate with the giant component of the network). For detection, we show that the probability that the detection time exceeds t is $\exp(-\Theta(t/\log t))$ in two dimensions, and $\exp(-\Theta(t))$ in three or more dimensions, under reasonable assumptions about the motion of the target. For percolation, we show that the probability that the percolation time exceeds t is $\exp(-\Omega(t^{\frac{d}{d+2}}))$ in all dimensions $d \geq 2$. We also give a sample application of this result by showing that the time required to broadcast a message through a mobile network with n nodes above the threshold density for existence of a giant component is $O(\log^{1+2/d} n)$ with high probability.

*Computer Science Division, University of California, Berkeley CA 94720-1776, U.S.A. Email: sinclair@cs.berkeley.edu. Supported in part by NSF grant CCF-0635153 and by a UC Berkeley Chancellor’s Professorship.

†Computer Science Division, University of California, Berkeley CA 94720-1776, U.S.A. Email: stauffer@cs.berkeley.edu. Supported by a Fulbright/CAPES scholarship and NSF grants CCF-0635153 and DMS-0528488.

1 Introduction

A principal focus in wireless network research today is on mobile ad hoc networks, in which nodes moving in space cooperate to relay packets on behalf of other nodes without any centralized infrastructure. Although the *static* properties of such networks are by now quite well understood mathematically, the additional challenges posed by node mobility have so far received relatively little attention from the theory community. In this paper we consider a mathematical model for mobile wireless networks, which we call *mobile geometric graphs* and which is a natural extension of the widely studied *random geometric graphs* model of static networks. We study two fundamental problems in this model: the *detection* problem (time until a fixed or moving target is detected by the network), and the *percolation* problem (time until a given node is able to communicate with many other nodes).

In the random geometric graph (RGG) model [27], nodes are distributed in a region $S \subseteq \mathbb{R}^d$ according to a Poisson point process of intensity λ (i.e., the number of nodes in any subregion $A \subseteq S$ is Poisson with mean $\lambda|A|$, where $|A|$ is the volume of A). Two nodes are connected by an edge iff their distance is at most r , where the parameter r is the *transmission range* that specifies the distance over which nodes may send and receive information; since the structure of the RGG depends only on the product $\lambda|B_r|$ (where B_r is the radius- r ball in \mathbb{R}^d) [6], we may fix r so that $|B_r| = 1$ and parameterize the model on λ only. We shall take S to be a cube of volume n/λ (so that the expected number of nodes in S is n), and consider the limiting behavior as $n \rightarrow \infty$.

Clearly, increasing λ increases the average degree of the nodes. As is well known, there are two critical values of λ at which the connectivity properties of the RGG undergo a significant change. First there is the *percolation threshold* $\lambda = \lambda_c$ (a constant that depends on the dimension d), so that if $\lambda > \lambda_c$ the network w.h.p.¹ has a unique “giant” component containing a constant fraction of the nodes, while if $\lambda < \lambda_c$ all components have size $O(\log n)$ w.h.p. [27]. Second, at the *connectivity threshold* $\lambda = \log n$, the network becomes connected w.h.p. [19]. The percolation threshold $\lambda = \lambda_c$ occurs also in the infinite-volume limit where $S = \mathbb{R}^d$, in which case the giant component is the unique infinite component (or “infinite cluster”) with probability 1. These and other fundamental properties of RGGs are extensively discussed in the book of Penrose [27]; see also [17] for additional results on thresholds. There are a host of theoretical results on routing and other algorithmic questions on static RGGs (see the Related Work section below for a partial list). Naturally, most of these consider networks above the connectivity threshold.

A central feature of many ad hoc networks is the fact that the nodes are moving in space. This is the case, for example, in vehicular networks (where sensors are attached to cars, buses or taxis), surveillance and disaster recovery applications where mobile sensors are used to survey an area, and pocket-switched networks based on mobile communication devices such as cellphones. Such networks are also frequently modeled using RGGs, augmented by motion of the nodes. We will employ the following model, which we refer to as *mobile geometric graphs* (MGGs) and which is essentially equivalent to the “dynamic boolean model” introduced in [5] in the context of dynamic continuum percolation. We begin at time 0 with a (infinite)² RGG G_0 in $S = \mathbb{R}^d$. Nodes move independently in continuous time according to Brownian motion with variance s^2 ; here s is a *range of motion* parameter, which we assume is constant to ensure a realistic model. We observe the nodes at discrete time steps (so the displacement of a node in each direction in each time step is normally distributed with mean 0 and variance s^2). It is not hard to verify that this produces a Markovian sequence of RGGs G_0, G_1, \dots , all with the same value of λ . Note that, while each G_i

¹We shall take the phrase “w.h.p.” (“with high probability”) to mean “with probability tending to 1 as $n \rightarrow \infty$.”

²Passing to infinite volume is a standard device that eliminates boundary effects in a finite region; with a little more technical effort the model and results can be extended to finite regions with a suitable convention—such as reflection or wraparound—to handle motion of nodes at the boundaries. See, e.g., Corollary 1.3 below.

is a RGG, there are correlations over time; it is this feature that makes mobility challenging to analyze.

Once mobility is injected, the questions of interest naturally change from those in the static case. For example, connectivity no longer plays such a central role because mobility may allow nodes u, v to exchange messages even in the absence of a path between them at any given time: namely, u can route its message to v along a time-dependent path, opportunistically using other nodes to relay the message towards v . Networks of this kind are often termed “delay tolerant networks” [14]. This allows us to focus not on the rather artificial connectivity regime (where λ grows with n), but instead on the case where λ (and hence the average degree) is constant. Keeping λ constant is obviously highly desirable as it makes the model much more realistic and scalable.

There are rather few rigorous results on wireless networks with mobile nodes, and those that do exist typically either make unrealistic assumptions about node mobility (such as unbounded range of motion [18, 11, 8] or no change in direction [25]), or work in the connectivity regime which, as we have seen, requires unbounded density or transmission range [10, 13]. (See the Related Work section for more details.) In this paper we study two fundamental questions for mobile networks assuming only constant average degree and bounded range of motion (i.e., constant values of the parameters λ and s).

Results

Detection. A central issue in surveillance and remote sensing applications is the ability of the network to detect a “target” $u \in S$ (which may be either fixed or moving), in the sense that there is a node within distance at most r of u . It is well known [26] that, for a static RGG, a fixed target can be detected only with constant probability unless the average degree grows with n . In the mobile case, we may hope to achieve detection *over time* with constant average degree ($\lambda = O(1)$, even *below* the percolation threshold). In this scenario, the detection time, T_{det} , is formulated as the number of steps until a target initially at the origin is detected by the MGG. Recent work of Liu et al. [25] shows that the detection time in two dimensions is exponentially distributed when the nodes of the network move in *fixed* directions. In the more realistic MGG model, we are able to prove the following result which holds in all dimensions (see Section 3):

Theorem 1.1. *In the MGG model with any fixed λ and range of motion $s > 0$, the detection time for a fixed target or a target moving under Brownian motion satisfies $\Pr[T_{\text{det}} \geq t] = \exp(-\Theta(t/\log t))$ for $d = 2$, and $\Pr[T_{\text{det}} \geq t] = \exp(-\Theta(t))$ for $d \geq 3$.*

The constants in the Θ here depend only on λ , s and the dimension d . Thus the tail of the detection time is exponential in three and higher dimensions, and exponential with a logarithmic correction in two dimensions. We note that, as is evident from the proof, this dichotomy between two and three dimensions reflects the difference between recurrent and transient random walks in \mathbb{Z}^2 and \mathbb{Z}^3 respectively. We also note that the upper bound in Theorem 1.1 holds for *arbitrary* motion of the target (provided it is independent of the motion of the nodes); and the lower bound holds for any “sufficiently random” motion of the target.

We should point out that, for the special case of a fixed target, a slightly stronger version of Theorem 1.1, with a tight constant in the exponent, follows from classical results on the “Wiener sausage” in continuum percolation. (This was pointed out to us by Yuval Peres [29]; see Related Work for details.) However, it is not clear how to extend this approach to the case of a moving target. Our proof is elementary and based on an application of the mass transport principle.

Percolation. A fundamental question in mobile networks is whether a node can efficiently communicate with other nodes even when the network is not connected at any given time. In the MGG model, this question may naturally be formulated by considering a constant intensity $\lambda > \lambda_c$ (i.e.,

above the percolation threshold) and asking how long it takes until a node initially at the origin belongs to the giant component (or the infinite component in the limit $n \rightarrow \infty$). We call this the *percolation time* T_{perc} . It should be clear that the percolation time can be used to derive bounds on other natural quantities, such as the time for a node to broadcast information to all other nodes (see Corollary 1.3 below). As far as we are aware the percolation time has not been investigated before, largely because previous work on RGGs has focused on networks above the connectivity threshold. However, it appears to be a fundamental question in the mobile context.

The detection time clearly provides a lower bound on the percolation time, so we may deduce from Theorem 1.1 above that $\Pr[T_{\text{perc}} \geq t]$ is at least $\exp(-O(t/\log t))$ for $d = 2$ and at least $\exp(-O(t))$ for $d \geq 3$. We are able to prove the following stretched exponential upper bound in all dimensions $d \geq 2$ (see Section 4):

Theorem 1.2. *In the MGG model with any fixed $\lambda > \lambda_c$ and range of motion $s > 0$, the percolation time for a node at the origin satisfies $\Pr[T_{\text{perc}} \geq t] = \exp(-\Omega(t^{d/(d+2)}))$ in all dimensions $d \geq 2$.*

Again, the constant in the Ω depends only on λ , s and d . There is a gap between this upper bound and the lower bound from Theorem 1.1. We conjecture that the true tail behavior of T_{perc} is $\exp(-\Theta(t/\log t))$ for $d = 2$ and $\exp(-\Theta(t))$ for $d \geq 3$.

Theorem 1.2 is the main technical contribution of the paper; we briefly mention some of the ideas used in the proof. The key technical challenge is the dependency of the RGGs G_i over time. To overcome this, we partition \mathbb{R}^d into subregions of suitable size and *couple* the evolution of the nodes in each subregion with those of a fresh Poisson point process of slightly smaller intensity $\lambda' < \lambda$ which is still larger than the critical value λ_c . After a number of steps Δ that depends on the size of the subregion, we are able to arrange that the coupled processes match up almost completely. As a result, we can conclude that our original MGG process, observed every Δ steps, contains a sequence of *independent* Poisson point processes with intensity $\lambda'' > \lambda_c$. (This fact, which we believe is of wider applicability, is formally stated in Proposition 4.1 in Section 4.) This independence is sufficient to complete the proof. The slack in the bound comes from the “delay” Δ .

To illustrate a sample application of Theorem 1.2, we consider the time taken to broadcast a message in a network of finite size. Consider a MGG in a cube of volume n/λ (so the expected³ number of nodes is n). Since the volume is finite, we need to modify the motion of the nodes to take account of boundary effects: following standard practice, we do this by turning the cube into a torus (so that nodes “wrap around” when they reach the boundaries). Suppose a message originates at an arbitrary node at time 0, and at each time step t each node that has already received the message broadcasts it to all nodes in the same connected component. (Here we are making the reasonable assumption that the speed of transmission is much faster than the motion of the nodes, so that messages can travel throughout a connected component before it is altered by the motion.) Let T_{bc} denote the time until all nodes have received the message.

Corollary 1.3. *In a MGG on the torus of volume n/λ with any fixed $\lambda > \lambda_c$ and range of motion $s > 0$, the broadcast time T_{bc} is $O(\log^{1+2/d} n)$ w.h.p. in any dimension $d \geq 2$.*

Related work

There are many theoretical results on routing and other algorithmic questions on (static) RGGs; we mention just a few highlights here. The seminal work of Gupta and Kumar [20, 21] (see also [15] for refinements) examined the information-theoretic *capacity* (or throughput) of such networks above the connectivity threshold, i.e., the number of bits per unit time that each node u can transmit

³The result can be adapted to the case of a *fixed* number of nodes n using standard “de-Poissonization” arguments [27]. See the Remark following the proof of Corollary 1.3 in Section 5.

to some (randomly chosen) destination node t_u in steady state, assuming constant size buffers in the network. The capacity per unit node is $\Theta(n^{-1/2})$, which tends to 0 as $n \rightarrow \infty$, suggesting a fundamental limitation on the scalability of such static networks.

The *detection* problem has received much attention. In the static case detection is essentially equivalent to *coverage* of the region S , which requires that the network be connected. In the absence of coverage, Balister et al. [3] determine the maximum diameter of the uncovered regions, while Dousse et al. [12] prove that, for any $\lambda > 0$, the detection time for a target moving in a fixed direction has an exponential tail. (Note that this is not a mobility result as the nodes are fixed.)

The question of *broadcasting* within the giant component of a RGG above the percolation threshold was recently analyzed by Bradonjić et al. [7], who also show that the graph distance between any two (sufficiently distant) nodes is at most a constant factor larger than their Euclidean distance. Cover times for random walks on (connected) RGGs were investigated by Avin and Ercal [1] and Cooper and Frieze [9], while the effect of physical obstacles that obstruct transmission was studied by Frieze et al. [16].

The scope of mathematically rigorous work on RGGs with mobility is much more limited. We briefly summarize it here.

Motivated by the fact mentioned above [20] that the capacity of static networks goes to zero as $n \rightarrow \infty$, Grossglauser and Tse [18] (see also [11]) showed how to exploit mobility to achieve constant capacity using a two-hop routing scheme. However, these results require the unrealistic assumption that nodes move a distance comparable to the diameter of the entire region S at each step. El Gamal et al. [13] study the tradeoff between capacity and delay in a realistic mobility model but above the connectivity threshold. Clementi et al. [8] show how to exploit mobility to enable broadcast in a RGG sufficiently far above the percolation threshold. However, this result again assumes that the range of motion of the nodes is unbounded (i.e., s grows with n).

As mentioned earlier, the detection problem was addressed by Liu et al. [25], assuming that each node moves continuously in a *fixed* randomly chosen direction; they show that the time it takes for the network to detect a target is exponentially distributed with expectation depending on the intensity λ . Also, for the special case of a stationary target, as observed in [23, 24] a slightly stronger version of Theorem 1.1, with tight constants in the exponent, can be deduced from classical results on continuum percolation: namely, in this case it is shown in [31] that (in continuous time) $\Pr[T_{\text{det}} \geq t] = \exp(-\lambda V_s(t))$, where $V_s(t)$ is the expected volume of the “Wiener sausage” of length $s^2 t$ (essentially the trajectory of a Brownian motion “fattened” by a disk of radius r). This volume in turn is known quite precisely [30, 4].

A model essentially equivalent to MGGs was introduced under the name “dynamic boolean model” by Van den Berg et al. [5], who studied the measure of the set of times at which an infinite component exists. Finally, recent work of Díaz et al. [10] in a similar model determines, for networks exactly at the connectivity threshold, the expected length of time for which the network stays connected (or disconnected) as the nodes move. However, this question makes sense only for very large values of λ (growing with n) and thus falls outside the scope of our investigations.

2 Preliminaries

For any $\ell \geq 0$, let B_ℓ be the d -dimensional ball centered at the origin with radius ℓ . Similarly, let Q_ℓ be the cube with side-length ℓ centered at the origin and with sides parallel to the axes of \mathbb{R}^d . For any point $z \in \mathbb{R}^d$ and set $A \subseteq \mathbb{R}^d$, we define $z + A$ as the Minkowski sum $z + A = \{y : y - z \in A\}$. The volume of a set $A \subset \mathbb{R}^d$ is denoted $|A|$.

Poisson point processes

A “point process” is a random collection of points in \mathbb{R}^d ; for a formal treatment of this topic, the

reader is referred to [31]. To avoid ambiguity, we refer to the points of a point process as *nodes* and reserve the word *points* for arbitrary locations in \mathbb{R}^d . We are mostly interested in *Poisson point processes*. A Poisson point process with intensity λ in a region $S \subseteq \mathbb{R}^d$ is defined by a single property: for every bounded Borel set $A \subseteq S$, the number of points in A is a Poisson random variable with mean $\lambda|A|$. We will make use of the following standard properties of Poisson point processes: (1) for disjoint sets $A, A' \subseteq S$, the numbers of points in A and in A' are independent; (2) conditioned on the number of nodes in A , each such node is located independently and uniformly at random in A ; (3) [thinning] if each node of a Poisson point process with intensity λ is deleted with probability p , the result is a Poisson point process with intensity $(1-p)\lambda$; (4) [superposition] the union of two Poisson point processes in S with intensities λ_1 and λ_2 is a Poisson point process with intensity $\lambda_1 + \lambda_2$. In some of our proofs we will make use of *non-homogeneous* Poisson point processes, whose intensity $\lambda(x)$ is a function of position $x \in \mathbb{R}^d$. In such a process the expected number of nodes in a set A is $\int_A \lambda(x)dx$.

Random geometric graphs

Fix parameters $\lambda, r \geq 0$, and let $S_n = Q_{(n/\lambda)^{1/d}}$ be the cube of volume n/λ in \mathbb{R}^d . Let Ξ_n be a Poisson point process over S_n with intensity λ . A random geometric graph (RGG) $\mathcal{G}(\Xi_n, r)$ is constructed by taking the node set to be the nodes of Ξ_n and creating an edge between every pair of nodes whose Euclidean distance is at most r . The parameter r is called the *transmission range*. Since Ξ_n is a Poisson point process, the expected number of nodes in $\mathcal{G}(\Xi_n, r)$ is n .

It is well known [6, 26] that as $n \rightarrow \infty$ the random graph model induced by $\mathcal{G}(\Xi_n, r)$ depends only on the product $\lambda|B_r|$. For this reason, we will always fix $r = r(d)$ so that $|B_r| = 1$ and parameterize the model only on λ . Note that with this convention, in the limit as $n \rightarrow \infty$, λ is also the expected degree of any node in $\mathcal{G}(\Xi_n, r)$.

Using a Poisson point process rather than a fixed number of nodes is a standard trick that simplifies the mathematics. Most results in this model can be translated to a model with a fixed number n of nodes in S_n using a technique known as “de-Poissonization” [27].

Many asymptotic properties of $\mathcal{G}(\Xi_n, r)$ as $n \rightarrow \infty$ are studied in the monograph by Penrose [27]. For example, it is known that $\lambda = \log n$ is a threshold for connectivity, in the sense that if $\lambda = \log n + \omega(1)$ then $\mathcal{G}(\Xi_n, r)$ is connected w.h.p., and if $\lambda = \log n - \omega(1)$ then $\mathcal{G}(\Xi_n, r)$ is disconnected w.h.p. Another important critical value is the *percolation threshold* $\lambda = \lambda_c$ (a constant that depends on the dimension d); if $\lambda > \lambda_c$ then w.h.p. $\mathcal{G}(\Xi_n, r)$ contains a unique “giant” connected component with $\Theta(n)$ nodes, while all other components are of size $O(\log^{d/(d-1)} n)$; on the other hand, if $\lambda < \lambda_c$ then w.h.p. all connected components of $\mathcal{G}(\Xi_n, r)$ have size $O(\log n)$. The value of λ_c is not known exactly in any dimension $d \geq 2$. However, for $d = 2$ the rigorous bounds $2.18 < \lambda_c < 10.60$ are known [26, Section 3.9], while Balister et al. [2] used Monte Carlo methods to deduce that $\lambda \in (4.508, 4.515)$ with confidence 99.99%.

Finally, we remark that in the limit as $n \rightarrow \infty$ (that is, when the Poisson point process is defined over the whole of \mathbb{R}^d) the percolation threshold $\lambda = \lambda_c$ still exists and is characterized by the appearance of a unique infinite component (or “infinite cluster”) with probability 1 for any $\lambda > \lambda_c$. In this limit the graph is disconnected with probability 1 for any value of λ .

Mobile geometric graphs

We define our mobile geometric graph (MGG) model by taking a Poisson point process with intensity λ in \mathbb{R}^d at time 0 and letting each node move in continuous time according to an independent Brownian motion. We sample the locations of the nodes at discrete time steps $i = 1, 2, \dots$ and use these locations to define a sequence of random geometric graphs with transmission range $r = r(d)$. We base our MGG model on the infinite volume \mathbb{R}^d to avoid having to handle boundary effects on the motion of the nodes. Results in this model can be translated to finite regions with a suitable convention—such as wraparound or reflection—to handle the motion of nodes at the boundaries.

More formally, let Π_0 be a Poisson point process with intensity λ over \mathbb{R}^d . We take a parameter $s \geq 0$, and with each node $v \in \Pi_0$ we associate an independent d -dimensional Brownian motion $\{W_v(i)\}_{i \geq 0}$ that starts at the location of v in Π_0 and has variance s^2 [22]. Now, for any $i \in \mathbb{Z}^+$, we define Π_i as the point process obtained by putting a node at $W_v(i)$ for each $v \in \Pi_0$. A MGG is then the collection of graphs $G = \{G_i\}_{i \geq 0}$ where $G_i = \mathcal{G}(\Pi_i, r)$ and $r = r(d)$ is fixed so that $|B_r| = 1$. (Note that, as in the static case, fixing the value of r may be done w.l.o.g.)

It is an easy consequence of the mass transport principle (see below) that each Π_i , viewed in isolation, is itself a Poisson point process with intensity λ . This means that the sequence $\{\Pi_i\}_{i \geq 0}$ is stationary and therefore, when viewed in isolation, G_i is a random geometric graph over \mathbb{R}^d . Thus, for example, if $\lambda > \lambda_c$ then each G_i contains an infinite component with probability 1.

Mass transport principle

For two points $x, y \in \mathbb{R}^d$ and a time step $i \geq 0$, we define $f_i(x, y)$ as the probability density function for a node located at position x at time 0 to be at position y at time i . Since nodes move according to d -dimensional Brownian motion, we have $f_i(x, y) = \frac{1}{(2\pi s^2 i)^{d/2}} \exp(-\frac{\|y-x\|_2^2}{2s^2 i})$.

In some situations, it is useful to regard $f_i(x, y)$ as a mass transport function. For example, suppose nodes are initially distributed according to a Poisson point process with intensity λ in a region $A \subseteq \mathbb{R}^d$; we may view this as a Poisson point process over \mathbb{R}^d with (non-homogeneous) intensity (or “mass function”) $\nu_0(x) = \lambda \mathbf{1}_{\{x \in A\}}$. Using the thinning and superposition properties, it is easy to check that the distribution of the nodes at time i is a Poisson point process with intensity $\nu_i(y) = \int_{\mathbb{R}^d} \nu_0(x) f_i(x, y) dx$. This interpretation can be used, for example, to show that in a MGG Π_i is a Poisson point process with intensity λ for all i , as claimed above.

3 Detection time

In this section we prove Theorem 1.1. We consider the detection time for a node u initially placed at the origin independently of the MGG G . We say that a node $v \in G$ *detects* u at time step i if the distance between u and v at time step i is at most r , and we define T_{det} as the first time that u is detected by some node of G . Our goal is to derive tight bounds for the tail of $\mathbf{Pr}[T_{\text{det}} \geq t]$. In the proof we consider the cases where u is either non-mobile or moves according to Brownian motion with variance s^2 . We discuss some extensions at the end of the section.

It will be convenient to restrict attention to the nodes of G that are initially inside the cube Q_L , where L is a suitably chosen parameter. We define $T_{\text{det}}(Q_L)$ as the first time a node initially inside Q_L detects u . Note that clearly $\mathbf{Pr}[T_{\text{det}}(Q_L) \geq t] \geq \mathbf{Pr}[T_{\text{det}} \geq t] = \lim_{L \rightarrow \infty} \mathbf{Pr}[T_{\text{det}}(Q_L) \geq t]$, where the limit exists since $\mathbf{Pr}[T_{\text{det}}(Q_L) \geq t]$ is monotone and bounded as a function of L . We let $X = (x_0, x_1, \dots, x_{t-1})$ be the locations of u in the first t steps. The following lemma relates the tail of $T_{\text{det}}(Q_L)$ to the tail of an analogous random variable for a single random node in Q_L .

Lemma 3.1. *We have $\mathbf{Pr}[T_{\text{det}}(Q_L) \geq t \mid X] = \exp(-\lambda L^d \mathbf{Pr}[\tau < t \mid X])$ and $\mathbf{Pr}[T_{\text{det}}(Q_L) \geq t] \geq \exp(-\lambda L^d \mathbf{Pr}[\tau < t])$, where τ is the first time that a node initially located u.a.r. in Q_L detects u .*

Proof. Let N be the number of nodes inside Q_L at time 0. Each of these nodes is initially located uniformly at random inside Q_L , and the motion of each node does not depend on the locations of the other nodes. If we fix a given value for X , then the first time that a given node v detects u does not depend on the other nodes of G and is distributed according to the conditional distribution of τ given X . (Note that this is not true if X is not fixed but random, because the random motion of u makes the relative displacements of the nodes of G with respect to u dependent.) Therefore, conditioning on N and X , we have $\mathbf{Pr}[T_{\text{det}}(Q_L) \geq t \mid N = n, X] = \mathbf{Pr}[\tau \geq t \mid X]^n$, which yields

$$\mathbf{Pr}[T_{\text{det}}(Q_L) \geq t \mid X] = \mathbf{E}_N \left[\mathbf{Pr}[\tau \geq t \mid X]^N \right] = \exp \left(-\lambda L^d \mathbf{Pr}[\tau < t \mid X] \right),$$

where we use the notation $\mathbf{E}_N[\cdot]$ to denote expectation with respect to the random variable N , and the last equality holds since N is Poisson with mean $\lambda|Q_L| = \lambda L^d$. For the lower bound, we appeal to Jensen's inequality to obtain

$$\begin{aligned}\mathbf{Pr}[T_{\text{det}}(Q_L) \geq t] &= \mathbf{E}_X[\mathbf{Pr}[T_{\text{det}}(Q_L) \geq t \mid X]] \\ &\geq \exp\left(-\lambda L^d \mathbf{E}_X[\mathbf{Pr}[\tau < t \mid X]]\right) \\ &= \exp\left(-\lambda L^d \mathbf{Pr}[\tau < t]\right),\end{aligned}$$

which completes the proof. \square

We now proceed to derive upper and lower bounds for $\mathbf{Pr}[\tau < t]$. Let v be a node initially located u.a.r. in Q_L . For time steps $i_1 \leq i_2$, let $M(i_1, i_2)$ be the expected number of time steps from i_1 to i_2 at which v detects u . We bound $M(0, t-1)$ as follows.

Lemma 3.2. *Let $t \in \mathbb{Z}^+$ and $X = (x_0, x_1, \dots, x_{t-1}) \in \mathbb{R}^{dt}$ be arbitrary. There exists a constant $c = c(d)$ and $L_0 = L_0(t, \max_i \|x_i\|_2)$ such that $ct/L^d \leq M(0, t-1) \leq t/L^d$ for all $L \geq L_0$.*

Proof. We use the mass transport principle. We assume the initial intensity $\nu_0(z) = \lambda \mathbf{1}_{\{z \in Q_L\}}$ and let $\nu_i(z)$ be the intensity at $z \in \mathbb{R}^d$ at time i , i.e., $\nu_i(z) = \int_{\mathbb{R}^d} \nu_0(y) f_i(y, z) dy$ for $i \geq 1$. At any time i , the probability that v detects u is given by the ratio between the amount of mass inside $x_i + B_r$ and the total amount of mass λL^d . Noting that for $i = 0$ this ratio is $\frac{\lambda|B_r|}{\lambda L^d} = 1/L^d$, we can write $M(0, t-1)$ as

$$\begin{aligned}M(0, t-1) &= \frac{1}{L^d} + \sum_{i=1}^{t-1} \frac{\int_{B_r} \nu_i(x_i + z) dz}{\lambda L^d} \\ &= \frac{1}{L^d} + \frac{1}{\lambda L^d} \sum_{i=1}^{t-1} \int_{B_r} \int_{Q_L} \lambda f_i(y, x_i + z) dy dz \\ &= \frac{1}{L^d} + \frac{1}{L^d} \sum_{i=1}^{t-1} \int_{B_r} \int_{x_i + z + Q_L} f_i(0, y') dy' dz,\end{aligned}\tag{1}$$

where the last step follows from the translation-invariance property of f_i . Since $|B_r| = 1$, we obtain the upper bound from $\int_{x_i + z + Q_L} f_i(0, y') dy' \leq \int_{\mathbb{R}^d} f_i(0, y') dy' = 1$.

For the lower bound, let $A_i = B_{s\sqrt{i+1}}$. We assume that L is sufficiently large such that $A_i \subseteq x_i + z + Q_L$. Then replacing the integral over $x_i + z + Q_L$ in (1) by an integral over A_i gives the lower bound

$$M(0, t-1) \geq \frac{1}{L^d} + \frac{1}{L^d} \sum_{i=1}^{t-1} \int_{B_r} \int_{A_i} f_i(0, y') dy' dz \geq \frac{1}{L^d} + \frac{1}{L^d} \sum_{i=1}^{t-1} \frac{|A_i||B_r|}{(2\pi s^2 i)^{d/2}} \exp(-1) \geq \frac{ct}{L^d},$$

where we used the fact that $\|y'\|_2 \leq s\sqrt{i+1}$ for all $y' \in A_i$, and the result holds for some constant $c = c(d)$. \square

Our goal is to write $M(i_1, i_2)$ conditioning on τ . Let $M'(y, i_1, i_2)$ be the expected number of time steps from i_1 to i_2 at which v detects u given that the relative location of v with respect to u at time $i_1 - 1$ is y . The next lemma gives lower and upper bounds for $M'(Y_i, i+1, i+t)$.

Lemma 3.3. *Let $i \in \mathbb{Z}^+$ be arbitrary. There exists an integer $t_0 = t_0(d, s)$ such that for all $t \geq t_0$ the following holds. There exist functions $m_1(t)$ and $m_2(t)$ such that $m_1(t) \leq M'(Y_i, i+1, i+t) \leq m_2(t)$ uniformly over $Y_i \in B_r$. Moreover, there are constants $c_1 = c_1(d)$ and $c_2 = c_2(d)$ such that*

$$m_1(t) \geq \begin{cases} c_1 \log t/s^d, & \text{for } d = 2 \\ c_1/s^d, & \text{for } d \geq 3 \end{cases} \quad m_2(t) \leq \begin{cases} c_2 \log t/s^d, & \text{for } d = 2 \\ c_2/s^d, & \text{for } d \geq 3 \end{cases}$$

The bounds for $m_1(t)$ hold both for the case where u does not move and for the case where u moves according to Brownian motion with variance s^2 . The bounds for $m_2(t)$ hold uniformly over X .

Proof. For any $j \in [1, t]$, let I_j be the indicator random variable for the event that v detects u at time $i+j$, assuming that at time i u is located at the origin and v is located at Y_i . Clearly, $M'(Y_i, i+1, i+t) = \sum_{j=1}^t \mathbf{E}[I_j]$. Recall that x_{i+j} is the location of u at time $i+j$. Hence,

$$\mathbf{E}[I_j] = \int_{B_r} \frac{1}{(2\pi s^2 j)^{d/2}} \exp\left(-\frac{\|x_{i+j} + z - Y_i\|_2^2}{2s^2 j}\right) dz \leq \frac{|B_r|}{(2\pi s^2 j)^{d/2}} = \frac{1}{(2\pi s^2 j)^{d/2}}.$$

The upper bound follows by setting $m_2(t) = \sum_{j=1}^t \frac{1}{(2\pi s^2 j)^{d/2}}$. Note that this upper bound holds for arbitrary X .

Now we derive the lower bound. We use the fact that $Y_i, z \in B_r$. If u is non-mobile, then $\|x_{i+j}\|_2 = \|x_i\|_2 = 0$ (recall that we assume x_i to be the origin) and from the triangle inequality we obtain $\|x_{i+j} + z - Y_i\|_2 \leq \|x_{i+j}\|_2 + \|z\|_2 + \|Y_i\|_2 \leq 2r$. Thus, $\mathbf{E}[I_j] \geq \frac{1}{(2\pi s^2 j)^{d/2}} \exp\left(-\frac{2r^2}{s^2 j}\right)$. We take j_0 to be the smallest integer such that $j_0 \geq r^2/s^2$, set $m_1(t) = \sum_{j=j_0}^t \frac{1}{(2\pi s^2 j)^{d/2}} \exp(-2)$, and the result follows since t is sufficiently large with respect to j_0 .

If u moves according to a Brownian motion with variance s^2 , we average over x_{i+j} to get

$$\mathbf{E}[I_j] = \int_{\mathbb{R}^d} \int_{B_r} \frac{1}{(2\pi s^2 j)^{d/2}} \exp\left(-\frac{\|x_{i+j} + z - Y_i\|_2^2}{2s^2 j}\right) \frac{1}{(2\pi s^2 j)^{d/2}} \exp\left(-\frac{\|x_{i+j}\|_2^2}{2s^2 j}\right) dz dx_{i+j}.$$

Let $a > 0$ be a constant and let $A_j = B_{as\sqrt{j}-2r}$. We set $a \geq 4$ so that $as\sqrt{j} - 2r \geq as\sqrt{j}/2$ for all $j \geq j_0 = \lceil r^2/s^2 \rceil$. We integrate over A_j instead of \mathbb{R}^d and then use the simple bounds $\|x_{i+j} + z - Y_i\|_2 \leq as\sqrt{j}$ and $\|x_{i+j}\|_2 \leq as\sqrt{j}$ for all $x_{i+j} \in A_j$ and $z, Y_i \in B_r$ to obtain

$$\mathbf{E}[I_j] \geq \frac{|A_j||B_r|}{(2\pi s^2 j)^d} \exp(-a^2) \geq \frac{a'}{(s^2 j)^{d/2}},$$

for some constant $a' = a'(d)$. Now, we set $m_1(t) = \sum_{j=j_0}^t \frac{a'}{(s^2 j)^{d/2}}$ and the result follows since t is sufficiently large with respect to j_0 . \square

Remark: The bounds for $M'(Y_i, i+1, i+t)$ change substantially from $d = 2$ to $d \geq 3$, reflecting the dichotomy between recurrent and transient random walks in \mathbb{Z}^2 and \mathbb{Z}^3 : v returns to a neighborhood of u infinitely often for $d = 2$ and only finitely often for $d \geq 3$. (Note that M' measures the expected number of returns of u to a neighborhood around v in a given time interval.)

We now use Lemma 3.3 to derive upper and lower bounds for $\mathbf{Pr}[\tau < t]$.

Lemma 3.4. *Let the functions m_1 and m_2 be as in Lemma 3.3. For any constant $\alpha > 0$ we have*

$$\frac{M(0, t-1)}{1 + m_2(t)} \leq \mathbf{Pr}[\tau < t] \leq \frac{M(0, (1+\alpha)t-1)}{1 + m_1(\alpha t)}.$$

Proof. We apply the straightforward equation

$$M(0, t-1) = \sum_{i=0}^{t-1} \mathbf{Pr}[\tau = i] (1 + \mathbf{E}_{Y_i} [M'(Y_i, i+1, t-1)]),$$

where the random variable Y_i denotes the relative location of v with respect to u given that $\tau = i$. Note that $Y_i \in B_r$ since the condition $\tau = i$ implies that the distance between v and u at time i is at most r . Using Lemma 3.3 we obtain

$$M(0, t-1) \leq \sum_{i=0}^{t-1} \mathbf{Pr}[\tau = i] (1 + \mathbf{E}_{Y_i} [M'(Y_i, i+1, i+t)]) \leq \mathbf{Pr}[\tau < t] (1 + m_2(t)). \quad (2)$$

Also, since $M(0, t-1)$ is non-decreasing with t , we can take an arbitrary constant $\alpha > 0$ and use the fact that $(1 + \alpha)t - 1 \geq i + \alpha t$ for all $i \in [0, t-1]$ together with Lemma 3.3 to write

$$\begin{aligned} M(0, (1 + \alpha)t - 1) &\geq \sum_{i=0}^{t-1} \mathbf{Pr}[\tau = i] (1 + \mathbf{E}_{Y_i} [M'(Y_i, i+1, (1 + \alpha)t - 1)]) \\ &\geq \sum_{i=0}^{t-1} \mathbf{Pr}[\tau = i] (1 + \mathbf{E}_{Y_i} [M'(Y_i, i+1, i + \alpha t)]) \\ &\geq \mathbf{Pr}[\tau < t] (1 + m_1(\alpha t)). \end{aligned} \quad (3)$$

The proof is completed by reorganizing (2) and (3). \square

We are now in position to conclude the proof of Theorem 1.1. Plugging Lemmas 3.2 and 3.3 into Lemma 3.4, and using Lemma 3.1, we obtain a constant $t_0 = t_0(d, s)$, and constants c_1, c_2, c_3, c_4 depending only on d , such that for all $t \geq t_0$ and sufficiently large L ,

$$\exp\left(-\frac{c_1 \lambda s^2 t}{s^2 + c_2 \log t}\right) \leq \mathbf{Pr}[T_{\det}(Q_L) \geq t] \leq \exp\left(-\frac{c_3 \lambda s^2 t}{s^2 + c_4 \log t}\right) \quad (4)$$

for $d = 2$, and

$$\exp\left(-c_1 \lambda s^d t\right) \leq \mathbf{Pr}[T_{\det}(Q_L) \geq t] \leq \exp\left(-c_3 \lambda s^d t\right) \quad (5)$$

for $d \geq 3$. Theorem 1.1 then follows by taking the limit as $L \rightarrow \infty$.

Remark: As should be clear from the proof, the upper bounds in (4) and (5) hold for arbitrary locations of u as long as u moves independently of the locations of the nodes of G . The lower bounds also hold in more generality: e.g., if u moves according to Brownian motion with variance $s'^2 \neq s^2$, or indeed with any motion that has sufficiently large “variance” in all directions. (Specifically, the lower bounds hold if the density $f'_i(\cdot, \cdot)$ for the motion of u after i steps satisfies the following property: there exist positive constants $a_1 = a_1(d)$ and $a_2 = a_2(d)$ such that $f'_i(x, x+z) \geq a_1/i^{d/2}$ for all $z \in B_{a_2\sqrt{i}}$, $x \in \mathbb{R}^d$, and $i \in \mathbb{Z}^+$.) On the other hand, adding a random drift to the nodes in G can change the detection time substantially: if each node v of G moves according to Brownian motion with drift μ_v and variance s^2 , where the μ_v are i.i.d. random variables, then our proof can be adapted to show that under mild conditions⁴ on the distribution of μ_v , $\mathbf{Pr}[T_{\det} \geq t] = \exp(-\Theta(t))$ in all dimensions $d \geq 2$ for arbitrary locations of u . We omit the details.

⁴Note that this statement cannot hold in full generality; if all nodes of G have the same drift, then this is equivalent to the case without drift up to translations of \mathbb{R}^d .

4 Percolation time

In this section we prove Theorem 1.2. We consider a MGG G with density $\lambda > \lambda_c$ (i.e., above the percolation threshold), and study the random variable T_{perc} defined as the first time at which a node u initially placed at the origin independently of the nodes of G belongs to the infinite component of G . We derive an upper bound for the tail $\mathbf{Pr}[T_{\text{perc}} \geq t]$ as $t \rightarrow \infty$.

We begin by stating a proposition that will be a key ingredient in our analysis. We consider a large cube $Q_K \subset \mathbb{R}^d$ and tessellate it into small cubes called “cells.” The proposition says that, if all cells have sufficiently many nodes at a given time i , then at time $i + \Delta$ for suitably large Δ the point process induced by the location of the nodes contains a fresh Poisson point process with only slightly reduced intensity inside a smaller cube $Q_{K'}$. We believe this result is of independent interest. With this in mind, we state the proposition below for a slightly more general setting than is needed here. Its proof is deferred to the end of the section.

Proposition 4.1. *Fix $K > \ell > 0$ and consider the cube Q_K tessellated into cells of side-length ℓ . Let Π_0 be an arbitrary point process at time 0 that contains at least $\beta\ell^d$ nodes in each cell of the tessellation for some $\beta > 0$. Let Π_Δ be the point process obtained at time Δ from Π_0 by allowing the nodes to move according to Brownian motion with variance s^2 . Fix $\epsilon \in (0, 1)$ and let Ξ be a fresh Poisson point process with intensity $(1 - \epsilon)\beta$. Then there exists a coupling of Ξ and Π_Δ and constants c_1, c_2, c_3 depending only on d such that, if $\Delta \geq \frac{c_1\ell^2}{s^2\epsilon^2}$ and $K' \leq K - c_2s\sqrt{\Delta \log \epsilon^{-1}} > 0$, the nodes of Ξ are a subset of the nodes of Π_Δ inside the cube $Q_{K'}$ with probability $1 - \frac{K^d}{\ell^d} \exp(-c_3\epsilon^2\beta\ell^d)$.*

Now we proceed to the proof of Theorem 1.2. We first take a sufficiently small parameter $\xi > 0$ such that $(1 - \xi)^2\lambda > \lambda_c$. (This is always possible as we are assuming $\lambda > \lambda_c$.) In what follows, we omit the dependencies of other parameters on λ and ξ as we are considering them to be fixed.

Let H_i be the event that u does *not* belong to the infinite component at time i . Then, the event $\{T_{\text{perc}} \geq t\}$ is equivalent to $\bigcap_{i=0}^{t-1} H_i$. We define an integer parameter $\Delta \geq 1$ and consider the process obtained by skipping every Δ time steps. (To simplify the notation we assume w.l.o.g. that t/Δ is an integer.) In other words, instead of looking at the event $\bigcap_{i=0}^{t-1} H_i$ we consider the event $\bigcap_{i=0}^{t/\Delta-1} H_{\Delta i}$, which we henceforth denote by \mathcal{H}_t . Since the occurrence of the event $\{T_{\text{perc}} \geq t\}$ implies \mathcal{H}_t we have $\mathbf{Pr}[T_{\text{perc}} \geq t] \leq \mathbf{Pr}[\mathcal{H}_t]$. Our goal in introducing Δ is to allow nodes to move further between consecutive time steps; we will choose the value of Δ later.

Let $C = C(d) \geq 1$ be a sufficiently large constant and fix $L = Ct(1 + s)$. We will confine our attention to the cube Q_{2L} . We take a parameter $\ell > 0$ and tessellate Q_{2L} into cubes of side-length ℓ (see Figure 1(a)). We refer to each such cube as a “cell.” Later we will tie together the values of ℓ and Δ , and will choose ℓ to optimize our upper bound for $\mathbf{Pr}[\mathcal{H}_t]$. For the moment we only assume that the tessellation is non-trivial in the sense that both ℓ and L/ℓ are $\omega(1)$ as functions of t .

For each time step i , the expected number of nodes inside a given cell is $\lambda\ell^d$. We say that a cell is *dense* at time i if it contains at least $(1 - \xi)\lambda\ell^d$ nodes, where ξ is as defined earlier. Let D_i be the event that *all* cells are dense at time i , and let $\mathcal{D}_t = \bigcap_{i=0}^{t/\Delta-1} D_i$. The lemma below shows that \mathcal{D}_t occurs with high probability.

Lemma 4.2. *With the above notation, $\mathbf{Pr}[\mathcal{D}_t] \geq 1 - \frac{tL^d}{\Delta\ell^d} \exp(-\xi^2\lambda\ell^d/2)$.*

Proof. At any given step i , by a standard large deviation bound for a Poisson r.v. (cf. Lemma A.1), a cell has more than $(1 - \xi)\lambda\ell^d$ nodes with probability at least $1 - \exp(-\xi^2\lambda\ell^d/2)$. The proof is completed by taking the union bound over all $(L/\ell)^d$ cells and t/Δ time steps. \square

Recall that x_i is the location of u at time i . Define E_i as the event that x_i is located inside $Q_{L/3}$ at time i , and let $\mathcal{E}_t = \bigcap_{i=0}^{t/\Delta-1} E_{\Delta i}$. The next lemma bounds the probability that u never leaves $Q_{L/3}$.




Figure 1: (a) The cubes Q_{2L} and Q_L and the tessellation of Q_{2L} into cubes of side-length ℓ . (b) The cube $Q_{L/3}$, the locations x_0 and x_i of node u at time steps 0 and i respectively, and the cube S_i .

Lemma 4.3. *There exists a constant $c = c(d)$ such that $\Pr[\mathcal{E}_t] \geq 1 - \exp(-ct)$.*

Proof. We fix a time step i and then apply the union bound over time. The event E_i corresponds to u not moving a distance more than $L/6$ in any dimension. Therefore,

$$\begin{aligned} \Pr[E_i] &= \int_{Q_{L/3}} f_i(0, x_i) dx_i \\ &= \left[1 - 2 \int_{L/6}^{\infty} \frac{1}{\sqrt{2\pi s^2 i}} \exp\left(-\frac{y^2}{2s^2 i}\right) dy \right]^d \geq 1 - 2d \int_{L/6}^{\infty} \frac{1}{\sqrt{2\pi s^2 i}} \exp\left(-\frac{y^2}{2s^2 i}\right) dy. \end{aligned}$$

Then we use a standard large deviation bound for the Normal distribution (see Lemma A.2) to conclude that

$$\Pr[E_i] \geq 1 - \frac{12ds\sqrt{i}}{\sqrt{2\pi}L} \exp\left(-\frac{L^2}{72s^2 i}\right).$$

Since the bound above decreases with i we can conclude that

$$\Pr[\mathcal{E}_t] \geq 1 - \frac{t}{\Delta} \frac{12ds\sqrt{t}}{\sqrt{2\pi}L} \exp\left(-\frac{L^2}{72s^2 t}\right),$$

and the result follows from $\Delta \geq 1$ and $L \geq st$. \square

For each time step i , we define S_i to be the cube $Q_{L/3}$ shifted randomly so that x_i (the location of u at time i) is uniformly random in S_i (see Figure 1(b)). A *crossing component* of S_i is a connected set of nodes within S_i that contains a path connecting every pair of opposite faces of S_i . (A path connects two faces of S_i if each face is within distance r of one of its endpoints.)

For each i , let K_i be the event that all the crossing components of S_i are contained in the infinite component at time i . (For definiteness we assume that K_i holds if S_i has no crossing component.) Let $\mathcal{K}_t = \bigcap_{i=0}^{t/\Delta-1} K_{\Delta i}$. The next lemma follows by a result of Penrose and Pisztora [28, Theorem 1].

Lemma 4.4. *For any $\lambda > \lambda_c$, there exists a constant $c = c(d)$ such that $\Pr[\mathcal{K}_t] \geq 1 - \exp(-ct^{d-1})$.*

Proof. By stationarity we know that $\Pr[K_i]$ is the same for all i . For any fixed i , [28, Theorem 1] gives that $\Pr[K_i] \geq 1 - \exp(-c'L^{d-1})$ for some constant c' . (In fact, [28, Theorem 1] handles an event more restrictive than K_i , which among other things considers *unique* crossing components of S_i .) Using the union bound we obtain $\Pr[\mathcal{K}_t] \geq 1 - (t/\Delta) \exp(-c'L^{d-1})$ and the result follows since $L \geq t$. \square

We now proceed to derive a bound on $\Pr[\mathcal{H}_t]$. We take H'_i to be the event that u does not belong to a crossing component of S_i at time i , and define $\mathcal{H}'_t = \bigcap_{i=0}^{t/\Delta-1} H'_{\Delta i}$. Note that H'_i is a *decreasing* event, in the sense that if H'_i occurs then it also occurs after removing any arbitrary collection of nodes from the MGG G . Clearly $\mathcal{H}_t \cap \mathcal{K}_t \subseteq \mathcal{H}'_t \cap \mathcal{K}_t$. By elementary probability,

$$\Pr[\mathcal{H}_t] \leq \Pr[\mathcal{H}'_t \cap \mathcal{D}_t \mid \mathcal{E}_t] + \Pr[\mathcal{D}_t^c] + \Pr[\mathcal{E}_t^c] + \Pr[\mathcal{K}_t^c]. \quad (6)$$

Note that we use \mathcal{K}_t only to replace \mathcal{H}_t by \mathcal{H}'_t in (6); this helps to control the dependencies among time steps, since \mathcal{H}'_t is an event restricted to the cubes S_i while \mathcal{H}_t is an event over the whole of \mathbb{R}^d . We use \mathcal{E}_t only to ensure that $S_i \subset Q_L$, which allows us to focus on the portion of G inside Q_L . Note that \mathcal{E}_t is independent of G , so this conditioning does not affect G .

Now we set $\Delta = \lceil C^2 \ell^2 / s^2 \rceil$, where C is the constant in the definition of L . The main step in our proof is the lemma below.

Lemma 4.5. *Let $\lambda > \lambda_c$ and $\xi > 0$ be such that $(1 - \xi)^2 \lambda > \lambda_c$. Let C be large enough in the definition of L and Δ . There exist constants $c = c(d)$ and $t_0 = t_0(d)$ such that, for all $t \geq t_0$, we have*

$$\Pr[\mathcal{H}'_t \cap \mathcal{D}_t \mid \mathcal{E}_t] \leq \exp(-ct/\Delta).$$

Proof. We start by writing

$$\Pr[\mathcal{H}'_t \cap \mathcal{D}_t \mid \mathcal{E}_t] \leq \prod_{i=0}^{t/\Delta-1} \Pr[H'_{\Delta i} \mid \mathcal{H}'_{\Delta(i-1)} \cap \mathcal{D}_{\Delta(i-1)} \cap \mathcal{E}_t]. \quad (7)$$

(Here, for notational convenience, we assume that $\mathcal{H}'_{-\Delta} \cap \mathcal{D}_{-\Delta} \cap \mathcal{E}_t = \mathcal{E}_t$.)

We now derive an upper bound for $\Pr[H'_{\Delta i} \mid \mathcal{H}'_{\Delta(i-1)} \cap \mathcal{D}_{\Delta(i-1)} \cap \mathcal{E}_t]$. We start with a high level overview of the proof. Let $\Xi_{\Delta(i-1)}$ be the (not necessarily Poisson) point process obtained from the nodes of $\Pi_{\Delta(i-1)}$ (the MGG at time $\Delta(i-1)$) under the condition $\mathcal{H}'_{\Delta(i-1)} \cap \mathcal{D}_{\Delta(i-1)}$. Note that $\Xi_{\Delta(i-1)}$ is conditioned only on events that occur between time 0 and time $\Delta(i-1)$; therefore, the motion of the nodes of $\Xi_{\Delta(i-1)}$ from time $\Delta(i-1)$ to Δi is independent of the condition. Since all cells are assumed dense at time $\Delta(i-1)$, using Proposition 4.1 we can construct an *independent* Poisson point process $\Xi'_{\Delta(i-1)}$ and couple it with $\Xi_{\Delta(i-1)}$ so that at time Δi the nodes of $\Xi'_{\Delta i}$ in Q_L are a subset of the nodes of $\Xi_{\Delta i}$. Moreover, we can ensure that $\Xi'_{\Delta i}$ has intensity larger than λ_c in Q_L , and thus conclude that u will belong to a crossing component of $S_{\Delta i}$ with constant probability. Using this, we can upper bound each term of the product in (7) by a constant strictly smaller than 1, which gives $\Pr[\mathcal{H}'_t \cap \mathcal{D}_t \mid \mathcal{E}_t] \leq \exp(-ct/\Delta)$ for some constant $c = c(d)$.

Turning now to the details, we can invoke Proposition 4.1 with $\beta = (1 - \xi)\lambda$, $\epsilon = \xi$, $K = 2L$, and $K' = L$ to obtain that, conditioned on $\mathcal{H}'_{\Delta(i-1)} \cap \mathcal{D}_{\Delta(i-1)}$, at time Δi the nodes of the MGG in Q_L contain a fresh Poisson point process with intensity $(1 - \xi)^2 \lambda > \lambda_c$ with probability at least $1 - \exp(-c'\xi^2 \lambda \ell^d)$, for some constant c' . Since $H'_{\Delta i}$ is a decreasing event⁵, if we define $H''_{\Delta i}$ as the event $H'_{\Delta i}$ restricted to the nodes of the fresh Poisson point process, then $H'_{\Delta i} \subseteq H''_{\Delta i}$ and

$$\Pr[H'_{\Delta i} \mid \mathcal{H}'_{\Delta(i-1)} \cap \mathcal{D}_{\Delta(i-1)} \cap \mathcal{E}_t] \leq \Pr[H''_{\Delta i}] + \exp(-c'\xi^2 \lambda \ell^d),$$

as $H''_{\Delta i}$ does not depend on the condition. Since the intensity of the fresh Poisson point process is $(1 - \xi)^2 \lambda > \lambda_c$, [28, Theorem 1] implies that, with probability $1 - \exp(-c''L^{d-1})$ for some constant c'' , a constant fraction of the volume of $S_{\Delta i}$ is within distance r of at least one node in

⁵Note that we defined $H'_{\Delta i}$ in terms of crossing components precisely in order to make it a decreasing event; otherwise, we could have defined it in terms of the largest connected component of $S_{\Delta i}$.

a crossing component of $S_{\Delta i}$. Since at time Δi , u is located uniformly at random inside $S_{\Delta i}$, u belongs to a crossing component of $S_{\Delta i}$ with probability at least $c''' > 0$, where c''' is a constant. Hence we have

$$\Pr[\mathcal{H}'_t \cap \mathcal{D}_t \mid \mathcal{E}_t] \leq \prod_{i=0}^{t/\Delta-1} \left[1 - c''' + \exp(-c'' L^{d-1}) + \exp(-c' \xi^2 \lambda \ell^d) \right].$$

Since L and ℓ go to infinity with t , for t sufficiently large each factor in the above product can be made strictly smaller than 1, which concludes the proof of Lemma 4.5. \square

Finally, we plug Lemmas 4.2–4.5 into (6) and obtain the following upper bound on $\Pr[\mathcal{H}_t]$:

$$\Pr[T_{\text{perc}} \geq t] \leq \Pr[\mathcal{H}_t] \leq \exp(-ct/\Delta) + \exp(-c\ell^d) + \exp(-ct) + \exp(-ct^{d-1}).$$

(Here $c = c(d)$ is a generic constant.) In order to minimize this upper bound we choose ℓ so that $\ell^d = \Theta(t/\Delta)$, which yields

$$\Pr[T_{\text{perc}} \geq t] \leq \exp\left(-ct^{\frac{d}{d+2}}\right)$$

for all sufficiently large t , where c is a constant depending on λ , s , and d . This completes the proof of Theorem 1.2. It remains to go back and prove Proposition 4.1.

Proof of Proposition 4.1

We will construct Ξ via three Poisson point processes. We start by defining Ξ'_0 as a Poisson point process over Q_K with intensity $(1 - \epsilon/2)\beta$. Recall that Π_0 has at least $\beta\ell^d$ nodes in each cell of Q_K . Then, in any fixed cell, Ξ'_0 has fewer nodes than Π_0 if Ξ'_0 has less than $\beta\ell^d$ nodes in that cell, which by a standard Chernoff bound (cf. Lemma A.1) occurs with probability larger than $1 - \exp\left(-\frac{\epsilon'^2(1-\epsilon/2)\beta\ell^d}{2}(1 - \epsilon'/3)\right)$ for ϵ' such that $(1 + \epsilon')(1 - \epsilon/2) = 1$. Since $\epsilon \in (0, 1)$ we have $\epsilon' \in (\epsilon/2, 1)$, and the probability above can be bounded below by $1 - \exp(-c\epsilon^2\beta\ell^d)$ for some constant $c = c(d)$. Let $\{\Xi'_0 \preceq \Pi_0\}$ be the event that Ξ'_0 has fewer nodes than Π_0 in every cell of Q_K . Using the union bound over cells we obtain

$$\Pr[\Xi'_0 \preceq \Pi_0] \geq 1 - \frac{K^d}{\ell^d} \exp(-c\epsilon^2\beta\ell^d). \quad (8)$$

If $\{\Xi'_0 \preceq \Pi_0\}$ holds, then we can map each node of Ξ'_0 to a unique node of Π_0 in the same cell. We will now show that we can couple the motion of the nodes in Ξ'_0 with the motion of their respective pairs in Π_0 so that the probability that an arbitrary pair is at the same location at time Δ is sufficiently large.

To describe the coupling, let v' be a node from Ξ'_0 located at $y' \in Q_K$, and let v be the pair of v' in Π_0 . Let y be the location of v in Q_K , and note that since v and v' belong to the same cell we have $\|y - y'\|_2 \leq \sqrt{d}\ell$. We will construct a function $g(z)$ that is smaller than the densities for the motions of v and v' to the location $y' + z$, uniformly for $z \in \mathbb{R}^d$. That is,

$$g(z) \leq \min\{f_{\Delta}(y', y' + z), f_{\Delta}(y, y' + z)\} = \frac{1}{(2\pi s^2 \Delta)^{d/2}} \exp\left(-\frac{\max\{\|z\|_2^2, \|y' + z - y\|_2^2\}}{2s^2 \Delta}\right) \quad (9)$$

for all $z \in \mathbb{R}^d$.

We set

$$g(z) = \frac{1}{(2\pi s^2 \Delta)^{d/2}} \exp\left(-\frac{(\|z\|_2 + \sqrt{d}\ell)^2}{2s^2 \Delta}\right). \quad (10)$$

Note that this definition satisfies (9) since by the triangle inequality $\|y' + z - y\|_2 \leq \|y' - y\|_2 + \|z\|_2$ and $\|y' - y\|_2 \leq \sqrt{d}\ell$. Define $\psi = 1 - \int_{\mathbb{R}^d} g(z)dz$. Then, with probability $1 - \psi$ we can use the density function $\frac{g(z)}{1-\psi}$ to sample a single location for the position of both v and v' at time Δ , and then set Ξ''_0 to be the Poisson point process with intensity $(1 - \psi)(1 - \epsilon/2)\beta$ obtained by *thinning* Ξ'_0 (i.e., deleting each node of Ξ'_0 with probability ψ). At this step we have crucially used the fact that the function $g(z)$ in (10) is oblivious of the location of v and, consequently, is independent of the point process Π_0 . (If one were to use the maximal coupling suggested by (9), then the thinning probability would depend on Π_0 , and Ξ''_0 would not be a Poisson point process.)

Let Ξ''_Δ be obtained from Ξ''_0 after the nodes have moved according to the density function $\frac{g(z)}{1-\psi}$. Thus we are assured that the nodes of the Poisson point process Ξ''_Δ are a subset of the nodes of Π_Δ and are independent of the nodes of Π_0 , where Π_Δ is obtained by letting the nodes of Π_0 move from time 0 to time Δ .

The next lemma shows that if Δ and $K - K'$ are large enough, then the integral of $g(z)$ inside the ball $B_{(K-K')/2}$ is larger than $1 - \epsilon/2$. (We are interested in the ball $B_{(K-K')/2}$ since for all $z \in Q_{K'}$ we have $z + B_{(K-K')/2} \subset Q_K$.)

Lemma 4.6. *If $\Delta \geq c \frac{\ell^2}{s^2 \epsilon^2}$ and $K - K' \geq c' s \sqrt{\Delta \log \epsilon^{-1}}$ for large enough c, c' , we may ensure that $\int_{B_{(K-K')/2}} g(z)dz \geq 1 - \epsilon/2$.*

Proof. Since $g(z)$ depends on z only via $\|z\|_2$, we integrate over $\rho = \|z\|_2$ and let $a\rho^{d-1}$ be the surface area of B_ρ , which gives

$$\int_{B_{(K-K')/2}} g(z)dz = \int_0^{(K-K')/2} \frac{a\rho^{d-1}}{(2\pi s^2 \Delta)^{d/2}} \exp\left(-\frac{(\rho + \sqrt{d}\ell)^2}{2s^2 \Delta}\right) d\rho.$$

Now we change variables to $\rho' = \frac{\rho + \sqrt{d}\ell}{s\sqrt{\Delta}}$ and set $\delta = \frac{\sqrt{d}\ell}{s\sqrt{\Delta}} = \epsilon\sqrt{d/c}$ and $K'' = \frac{K-K'}{2s\sqrt{\Delta}} + \delta$ to obtain

$$\int_{B_{(K-K')/2}} g(z)dz = \frac{a}{(2\pi)^{d/2}} \int_\delta^{K''} (\rho' - \delta)^{d-1} \exp(-\rho'^2/2) d\rho'. \quad (11)$$

Let $h(\delta) = \int_\delta^\infty (\rho' - \delta)^{d-1} \exp(-\rho'^2/2) d\rho'$. We will apply the Taylor expansion of $h(\delta)$ around $\delta = 0$. Note that $\frac{a}{(2\pi)^{d/2}} h(0) = 1$, and the derivative of $h(\delta)$ is

$$h'(\delta) = -(d-1) \int_\delta^\infty (\rho' - \delta)^{d-2} \exp(-\rho'^2/2) d\rho'.$$

In particular, the derivative increases with δ , and by Taylor's theorem, $h(\delta) \geq h(0) + \delta h'(0)$. Therefore, we have

$$\frac{a}{(2\pi)^{d/2}} h(\delta) \geq \frac{a}{(2\pi)^{d/2}} (h(0) + \delta h'(0)) = 1 - \delta \frac{ah'(0)}{(2\pi)^{d/2}}.$$

Note that $\delta \frac{ah'(0)}{(2\pi)^{d/2}}$ depends on the dimension only and can, for example, be made smaller than $\epsilon/4$ for sufficiently large c . Using equation (11) we get

$$\int_{B_{(K-K')/2}} g(z)dz \geq 1 - \epsilon/4 - \frac{a}{(2\pi)^{d/2}} \int_{K''}^\infty (\rho' - \delta)^{d-1} \exp(-\rho'^2/2) d\rho'. \quad (12)$$

Now note that $\rho'^{d-1} \exp(-\rho'^2/2) \leq c'' \exp(-\rho'^2/3)$ uniformly for $\rho' \in [0, \infty)$, where c'' is a constant depending only on d . Thus we have

$$\int_{K''}^{\infty} \frac{a}{(2\pi)^{d/2}} (\rho' - \delta)^{d-1} \exp(-\rho'^2/2) d\rho' \leq \frac{c'' a}{(2\pi)^{d/2}} \int_{K''}^{\infty} \exp(-\rho'^2/3) d\rho' \leq c''' \exp(-K''^2/3),$$

for a constant c''' depending only on d . For sufficiently large c' (and thus K''), the right hand side in (13) can be made smaller than $\epsilon/4$. Plugging this into (12) yields the lemma. \square

When $\{\Xi'_0 \preceq \Pi_0\}$ holds, Ξ''_{Δ} consists of a subset of the nodes of Π_{Δ} . Note that Ξ''_{Δ} is a *non-homogeneous* Poisson point process over Q_K . It remains to show that the intensity of Ξ''_{Δ} is strictly larger than $(1-\epsilon)\beta$ in $Q_{K'}$ so that Ξ can be obtained from Ξ''_{Δ} via thinning; since Ξ''_{Δ} is independent of Π_0 , so is Ξ .

For $z \in \mathbb{R}^d$, let $\mu(z)$ be the intensity of Ξ''_{Δ} . Since Ξ''_0 has no node outside Q_K , we obtain for any $z \in Q_{K'}$,

$$\mu(z) \geq (1-\psi)(1-\epsilon/2)\beta \int_{z+B_{(K-K')/2}} \frac{g(z-x)}{1-\psi} dx = (1-\epsilon/2)\beta \int_{B_{(K-K')/2}} g(x) dx,$$

where the inequality follows since $z + B_{(K-K')/2} \subset Q_K$ for all $z \in Q_{K'}$. From Lemma 4.6, we have $\int_{B_{(K-K')/2}} g(x) dx \geq 1 - \epsilon/2$. We then obtain $\mu(z) \geq (1-\epsilon/2)^2\beta \geq (1-\epsilon)\beta$, which is the intensity of Ξ . Therefore, when $\{\Xi'_0 \preceq \Pi_0\}$ holds, which occurs with probability given by (8), the nodes of Ξ are a subset of the nodes of Π_{Δ} , which completes the proof of Proposition 4.1.

5 Broadcast time

In this section we use Theorem 1.2 to prove Corollary 1.3 for a finite mobile network of volume n/λ .

We may relate the MGG model on the torus to a model on \mathbb{R}^d as follows. Let S_n denote the cube $Q_{(n/\lambda)^{1/d}}$. The initial distribution of the nodes is a Poisson point process over \mathbb{R}^d with intensity λ on S_n and zero elsewhere. We allow the nodes to move according to Brownian motion over \mathbb{R}^d as usual, and at each time step we project the location of each node onto S_n in the obvious way.

Now let $t = C \log^{1+2/d} n$ for some sufficiently large constant $C = C(d)$. The proof proceeds in three stages. First, we show that for any fixed $i \in [0, t-1]$, the giant component of G_i has at least one node in common with the giant component of G_{i+1} . This means that, once the message has reached the giant component, it will reach any node v as soon as v itself belongs to the giant component. Thus we can bound the broadcast time by (twice) the time until all nodes of G have been in the giant component.

In order to prove the above claim, let $\epsilon > 0$ be sufficiently small so that $(1-\epsilon)\lambda > \lambda_c$. We use the thinning property to split Π_i into two Poisson point processes, Π'_i and Π''_i , with intensities $(1-\epsilon)\lambda$ and $\epsilon\lambda$ respectively. Let G'_i and G''_{i+1} be the RGGs induced by Π'_i and Π''_{i+1} respectively. Then with probability $1 - e^{-\Theta(n^{1-1/d})}$ both G'_i and G''_{i+1} contain a giant component [28]. We show that at least one node from Π''_i belongs to both giant components. For any node v of Π''_i , the probability that v belongs to the giant component of G'_i is larger than some constant $c = c(d)$. Moreover, using the FKG inequality we can show that v belongs to the giant components of both G'_i and G''_{i+1} with probability larger than c^2 . Therefore, using the thinning property again, we can show that the nodes from Π''_i that belong to the giant components of both G'_i and G''_{i+1} form a Poisson point process with intensity $\epsilon\lambda c^2$, since c does not depend on Π''_i . Hence, there will be at least one such node inside S_n with probability $1 - e^{-\epsilon c^2 n}$, and this stage is concluded by taking the union bound over time steps i .

Our ultimate goal is to show that if $T_{\text{perc}} \leq t$ then all nodes of G receive the message being broadcast within $2t$ steps w.h.p. We proceed to the second stage of the proof, and show that the tail bound on T_{perc} from Theorem 1.2 also holds when applied to the finite region S_n defined above. Note that all the derivations in the proof of Theorem 1.2 were restricted to the cube Q_{2L} , where L was defined near the beginning of Section 4. Therefore, it is enough to show that Q_{2L} is contained inside S_n (so that the toroidal boundary conditions do not affect the result). But this holds for all sufficiently large n since $L = O(t) = O(\log^{1+2/d} n)$ while S_n has side-length $(n/\lambda)^{1/d}$.

The last stage of the proof consists of showing that adding a node u at the origin and calculating its percolation time (as we did in Theorem 1.2) is equivalent to calculating the percolation time of an arbitrary node of G . Note that, by a Chernoff bound, G has at most $(1+\delta)n$ nodes with probability larger than $1 - e^{-\Omega(n)}$ for any fixed $\delta > 0$. These nodes are indistinguishable, so letting ρ be the probability that an arbitrary node has percolation time at least t , we can use the union bound to deduce that this applies to *at least one* node in G with probability at most $(1+\delta)n\rho$. Let v be an arbitrary node. In order to relate ρ to the result of Theorem 1.2, we can use translation invariance and assume that v is at the origin. Then, by the ‘‘Palm theory’’ of Poisson point processes [31], ρ is equivalent to the tail of the percolation time for a node added at the origin, which is precisely $\Pr[T_{\text{perc}} \geq t]$. Thus finally, using Theorem 1.2 we get $\rho \leq \exp(-ct^{\frac{d}{d+2}})$, which can be made $o(1/n)$ by setting C sufficiently large. This completes the proof of Corollary 1.3.

Remark: It is easy to see that the above result also holds in the case where the MGG has *exactly* n nodes. The proof above shows that, by setting C large enough, we can ensure $\Pr[T_{\text{bc}} \geq t] = o(1/n)$ for the given value of t . Also, it is well known that a Poisson random variable with mean n takes the value n with probability $p = \Theta(1/\sqrt{n})$. Therefore, for a MGG with exactly n nodes, we have $\Pr[T_{\text{bc}} < t] = \frac{p - o(1/n)}{p} = 1 - o(1/\sqrt{n})$.

6 Some open questions

We hope our work will promote further mathematical research on mobile networks. Some natural open questions related to our results include the following:

1. What is the tight asymptotic behavior of the tail of the percolation time T_{perc} ? We conjecture that the upper bound of Theorem 1.2 can be tightened to match the lower bounds provided by $\Pr[T_{\text{det}} \geq t]$ in Theorem 1.1.
2. Let the *coverage time* T_{cov} be the time until *all* points of a finite region S are detected by the MGG. How does T_{cov} behave?
3. Let T_{comm} denote the time until a specific node u is able to send a message to a target node v , assuming that the other nodes cooperate maximally to achieve this following the broadcast protocol of Section 5. Since T_{comm} is bounded above by the time until both u and v simultaneously belong to the giant component, following along the lines of the proof of Theorem 1.2 for the tail of T_{perc} we can show that $\Pr[T_{\text{comm}} \geq t] = \exp(-\Omega(t^{\frac{d}{d+2}}))$. Can one show a substantially better upper bound on the tail of T_{comm} ?
4. What is the broadest (realistic) class of mobility models (i.e., beyond the Brownian motion we consider in this paper) for which our results still apply?
5. In our result on the broadcast time (Corollary 1.3), we assumed that messages can travel instantaneously throughout each connected component, which is reasonable in many applications (where, e.g., transmission through the air is effectively instantaneous in comparison to the motion of the nodes). How is this result affected by the assumption that messages may only travel a limited number of hops at each time step?

Acknowledgments

We thank Yuval Peres for helpful input on continuum percolation, and in particular for pointing out to us the connection between the detection problem and the Wiener sausage. We also thank David Tse for useful discussions on mobile wireless networks.

References

- [1] C. AVIN and G. ERCAL. On the cover time and mixing time of random geometric graphs. *Theoretical Computer Science* **380** (2007), pp. 2–22.
- [2] P. BALISTER, B. BOLLOBÁS and M. WALTERS. Continuum percolation with steps in the square or the disc. *Random Structures and Algorithms* **26** (2005), pp. 392–403.
- [3] P. BALISTER, Z. ZHENG, S. KUMAR and P. SINHA. Trap coverage: Allowing coverage holes of bounded diameter in wireless sensor networks. *Proceedings of the 28th IEEE Conference on Computer Communications*, 2009, pp. 19–25.
- [4] A.M. BEREZHOVSKII, YU.A. MAKHOVSKII and R.A. SURIS. Wiener sausage volume moments. *Journal of Statistical Physics* **57** (1989), pp. 333–346.
- [5] J. VAN DEN BERG, R. MEESTER and D.G. WHITE. Dynamic boolean model. *Stochastic Processes and their Applications* **69** (1997), pp. 247–257.
- [6] B. BOLLOBÁS and O. RIORDAN. *Percolation*. Cambridge University Press, 2006.
- [7] M. BRADONJIĆ, R. ELSÄSSER, T. FRIEDRICH, T. SAUERWALD and A. STAUFFER. Efficient broadcast on random geometric graphs. *Proceedings of the 21st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2010, pp. 1412–1421.
- [8] A. CLEMENTI, F. PASQUALE and R. SILVESTRI. MANETS: high mobility can make up for low transmission power. *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, 2009.
- [9] C. COOPER and A. FRIEZE. The cover time of random geometric graphs. *Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2009, pp. 48–57.
- [10] J. DÍAZ, D. MITSCHKE and X. PÉREZ-GIMÉNEZ. On the connectivity of dynamic random geometric graphs. *Proceedings of the 19th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2008, pp. 601–610.
- [11] S. DIGGAVI, M. GROSSGLAUSER and D. TSE. Even one-dimensional mobility increases adhoc wireless capacity. *IEEE Transactions on Information Theory* **51** (2005), pp. 3947–3954.
- [12] O. DOUSSE, C. TAVOULARIS and P. THIRAN. Delay of intrusion detection in wireless sensor networks. *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking (MobiCom)*, 2006, pp. 155–165.
- [13] A. EL GAMAL, J. MAMMEN, B. PRABHAKAR and D. SHAH. Throughput-delay trade-off in wireless networks. *Proceedings of the 23rd IEEE Conference on Computer Communications*, 2004, pp. 464–475.
- [14] K. FALL. A delay-tolerant network architecture for challenged internets. *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures and Protocols for Computer Communications*, 2003, pp. 27–34.
- [15] M. FRANCESCHETTI, O. DOUSSE, D. TSE and P. THIRAN. Closing the gap in the capacity of random wireless networks via percolation theory. *IEEE Transactions on Information Theory* **53** (2007), pp. 1009–1018.

- [16] A. FRIEZE, J. KLEINBERG, R. RAVI and W. DEBANY. Line-of-sight networks. *Combinatorics, Probability and Computing* **18** (2009), pp. 145–163.
- [17] A. GOEL, S. RAI and B. KRISHNAMACHARI. Sharp thresholds for monotone properties in random geometric graphs. *Proceedings of the 36th ACM Symposium on Theory of Computing (STOC)*, 2004, pp. 580–586.
- [18] M. GROSSGLAUSER and D. TSE. Mobility increases the capacity of ad hoc wireless networks. *IEEE Transactions on Networking* **10** (2002), pp. 477–486.
- [19] P. GUPTA and P.R. KUMAR. Critical power for asymptotic connectivity in wireless networks. In *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of W.H. Fleming*, W.M. McEneaney, G. Yin and Q. Zhang (eds.), Birkhäuser, Boston, 1998, pp. 547–566.
- [20] P. GUPTA and P.R. KUMAR. The capacity of wireless networks. *IEEE Transactions on Information Theory* **46** (2000), pp. 388–404. Correction in *IEEE Transactions on Information Theory* **49** (2000), p. 3117.
- [21] P. GUPTA and P.R. KUMAR. Internets in the Sky: The capacity of three dimensional wireless networks. *Communications in Information and Systems* **1** (2001), pp. 33–49.
- [22] I. KARATZAS and S.E. SHREVE. *Brownian Motion and Stochastic Calculus* (2nd ed.). Springer, 1991.
- [23] G. KESIDIS, T. KONSTANTOPOULOS and S. PHOHA. Surveillance coverage of sensor networks under a random mobility strategy. *Proceedings of the 2nd IEEE International Conference on Sensors*, 2003.
- [24] T. KONSTANTOPOULOS. Response to Prof. Baccelli’s lecture on Modelling of Wireless Communication Networks by Stochastic Geometry. *Computer Journal Advance Access*, 2009.
- [25] B. LIU, P. BRASS, O. DOUSSE, P. NAIN and D. TOWSLEY. Mobility improves coverage of sensor networks. *Proceedings of the 6th ACM International Conference on Mobile Computing and Networking (MobiCom)*, 2005.
- [26] R. MEESTER and R. ROY. *Continuum Percolation*. Cambridge University Press, 1996.
- [27] M. PENROSE. *Random Geometric Graphs*. Oxford University Press, 2003.
- [28] M. PENROSE and A. PISZTORA. Large deviations for discrete and continuous percolation. *Advances in Applied Probability* **28** (1996), pp. 29–52.
- [29] Y. PERES. Personal communication, February 2010.
- [30] F. SPITZER. Electrostatic capacity, heat flow, and Brownian motion. *Z. Wahrscheinlichkeitstheorie verw. Geb.* **3** (1964), pp. 110–121.
- [31] D. STOYAN, W.S. KENDALL and J. MECKE. *Stochastic Geometry and its Applications*. John Wiley & Sons, 2nd ed., 1995.

A Standard large deviation results

We use the following standard Chernoff bounds and large deviation results.

Lemma A.1 (Chernoff bound for Poisson). *Let P be a Poisson random variable with mean λ . Then, for any $0 < \epsilon < 1$,*

$$\Pr [P \geq (1 + \epsilon)\lambda] \leq \exp \left(-\frac{\lambda\epsilon^2}{2}(1 - \epsilon/3) \right),$$

and

$$\Pr [P \leq (1 - \epsilon)\lambda] \leq \exp \left(-\frac{\lambda\epsilon^2}{2} \right).$$

Lemma A.2 (Large deviation for normal). *Let N be a Normal random variable with mean 0 and variance σ^2 . Then, for any $x \geq 0$,*

$$\Pr [N \geq x] \leq \frac{\sigma}{\sqrt{2\pi}x} \exp \left(-\frac{x^2}{2\sigma^2} \right).$$